

NO-A166 428


DYNAMIC MULTI-ATTRIBUTE PERSONNEL ASSIGNMENT(U) TEXAS
UNIV AT AUSTIN CENTER FOR CYBERNETIC STUDIES A CHARNES
31 MAR 86 N00014-81-C-0236


1/2

UNCLASSIFIED

F/G 5/1

NL





MICROCOPY RESOLUTION TEST CHART

1

DYNAMIC MULTI-ATTRIBUTE
PERSONNEL ASSIGNMENT

(ONR Contract N00014-81-C-0236)

Final Technical Report

by

A. Charnes


AD-A166 428

**CENTER FOR
CYBERNETIC
STUDIES**

The University of Texas
Austin, Texas 78712

DTIC
ELECTE
APR 08 1986
S D

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited



4
3 7
9

1

**DYNAMIC MULTI-ATTRIBUTE
PERSONNEL ASSIGNMENT**

(ONR Contract N00014-81-C-0236)

Final Technical Report

by

A. Charnes

March 31, 1986

**DTIC
ELECTE
APR 08 1986
S D D**

This report was distributed to the Defense Technical Information Center
on April 3, 1986.

CENTER FOR CYBERNETIC STUDIES

**A. Charnes, Director
College of Business Administration 5.202
The University of Texas at Austin
Austin, Texas 78712-1177
(512) 471-1821**

DISTRIBUTION STATEMENT A

**Approved for public release
Distribution Unlimited**

FINAL REPORT

ONR Contract N00014-81-C-0236

The work of this contract encompassed significant new developments in network theory and in computational methods to the extent that a system for monitoring performance and for evaluation of policy with regard to Sea-Shore Rotation was produced which was operational in real time. All software developments, documentation, and deliverables were delivered long before the concluding date of the contract. Details of the system and a manual for its usage are included ~~here~~ as Appendices 1 and 2. As with other previous efforts the progress of this was repeatedly delayed by organizational changes and reassignments of the responsible Navy officers and with new changes in directions which conflicted with previous assignments. The closure, rather than extension, of this project means that a significant capability in terms of the research team which developed the system is lost to the Navy and that costly majorally duplicating new extensive effort will be required to achieve the operational capabilities which would have been available from an inexpensive, modest extension of this contract. It must be recognized, of course, that the decision not to extend must have been based on higher priorities for the Navy, at this time.



Accession For	
NTIS	<input checked="checked" type="checkbox"/>
CRA&I	<input checked="checked" type="checkbox"/>
DTIC	<input type="checkbox"/>
TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

APPENDIX 1

Research Report CCS 495

**The GPSSR System to Support
Management of Policy and Execution of
The U.S. Navy's Sea-Shore Rotation Program**

by

A. Charnes
W.W. Cooper***
B. Golany

V. Lovegren *
W.T. Mayfield**
M. Wolfe

September 1984

Revised November 1984

* University of Illinois at Chicago Circle

** United States Navy

This research was partly supported by ONR Contract N00014-81-C-0236 and ONR Contract N00014-81-K-0295 with the Center for Cybernetic Studies, The University of Texas at Austin and part of (***) by ONR Contract N00014-82-C-0411 with the School of Urban and Public Affairs at Carnegie-Mellon University. Reproduction whole or in part is permitted for any purpose of the United States Government.

CENTER FOR CYBERNETIC STUDIES

A. Charnes, Director
College of Business Administration 5.202
The University of Texas at Austin
Austin, Texas 78712-1177
(512) 471-1821

ABSTRACT

This paper is an exposition of the GPSSR system to support management of policy and execution of the U.S. Navy's Enlisted Personnel Sea Shore Rotation Program. Its components include (1) a new model of constrained network goal programming type; (2) newly developed algorithms for use with models of this class; (3) computer software and informatics developed to implement these algorithms, plus the software and informatics for other modules of the system including (4) decision support tools for report generation and monitoring capabilities.

Key Words:

Goal Programming
Sea Shore Rotation
Valve Arc Method
Constrained Networks

CONTENTS

1.0 INTRODUCTION	3
2.0 BACKGROUND	6
3.0 Model and System Overview	11
4.0 GPSSR System Modules	14
4.1 Extraction and Separation of Data	14
4.2 Calculation Of Transition Rates	15
4.2.1 Obtaining Transition Totals	15
4.2.2 Calculating The Smoothed Rates	16
4.2.3 User Interaction	17
4.3 Network Generator	17
4.3.1 Generating Bounds	18
4.3.2 Attaching Goals and Costs	20
4.4 Network Optimizer (VICNET)	21
4.5 Report Generator	22
5.0 Monitoring Features	32

5.1	Displaying Statistics from the EMR	32
5.2	Additional Monitoring Capabilities	39
6.0	Work In Progress	41
6.1	Developments In The Theory	41
6.2	Technical Enhancements/Refinements	42
7.0	Summary and Conclusions	43
7.1	Summary	43
7.2	Conclusions and Directions of Further Work	43
APPENDIX A		A1
APPENDIX B		B1

1.0 INTRODUCTION

This report presents the results of research by the Center for Cybernetic Studies to provide a system that will support the management of policy and execution of the Navy's Enlisted Personnel Sea Shore Rotation Program. This system consists of several integrated components each of which represents an advance in the present state of modeling and computerized algorithms. These components include (1) a new model of constrained network goal programming type; (2) newly developed algorithms for use with models of this class; (3) computer software and informatics developed to implement these algorithms, plus the software and informatics for other modules of the system including (4) decision support tools for report generation and monitoring capabilities.

The system was developed with participation by the staff of the Navy Military Personnel Command (NMPC) ¹ and the Manpower and Career Planning Research Group at Carnegie-Mellon University, and is designed to meet present Navy requirements for both planning and policy evaluations.

¹ Special thanks are due to CDR E. L. Kainer who provided guidance and help at many points in the course of these developments and to ETCM William F. Hinkel who helped in collection of the data, formulation of the model, and interpretation of our results. LT. Gareth Habel was also helpful at critical junctures in the developments covered in this report.

Because the model utilizes a goal programming approach, we refer to it as the Goal Programming Sea Shore Rotation (GPSSR) model.

The GPSSR model is designed for use in planning and scheduling personnel flows and for evaluating the consequences of such flows relative to Navy goals and policies. In principle, it examines all possible personnel flows and selects the ones that come closest to meeting all goals while honoring the specified policy and operational constraints. It also has the capability of evaluating alternatives in policy or operational constraints in terms of their effects on goal achievements. Thus, on the one hand, it shows what is required to do the best possible job under the given constraints and, on the other hand, it allows the exploration of alternatives which the user might wish to consider. By providing a consistent basis for both policy and operational planning through its decision support tools, GPSSR also provides a framework for policy execution monitoring.

Special algorithms developed at the Center for Cybernetic Studies which have now been incorporated in computer software, make it possible to provide the above capabilities efficiently and effectively. Solution times of a minute or less for a complete detailing community have already been achieved and further reductions in these times are possible. This provides Navy managers with a capability for policy analysis and planning at various desired levels of detail without accompanying delays: a Navy manager may try a variety of personnel rotation flow alternatives and have

the consequences immediately available for consideration. The report generation and graphic display capabilities contained in the system further aid the manager's decision making.

2.0 BACKGROUND

The management of Navy enlisted personnel includes the continuous task of planning for and executing sea shore rotation policies. This management task is described in the Navy's Enlisted Transfer Manual NAVPERS 15909C Articles 3.0-3.01 as follows:

"The system for the planned reassignment of personnel among the various types of duty is designed to

- Promote maximum readiness and stability both afloat and ashore.
- Permit equitable opportunity for personnel to serve in duty they consider desirable.

Rotation among sea, shore, and overseas activities is directly influenced by the number of personnel available for assignment, billets authorized, PCS funds, and qualifications of the individual."

Deciding upon rotation policies which satisfy a variety of oftentimes conflicting objectives is a large and complex problem with many different dimensions. Even when restricted to enlisted personnel, each of more than 250 detailing communities must be individually considered, and, for effective planning, qualifications like the following are involved:

- Paygrade
- Rating
- Subspecialty (NEC) Community
- Obligated Service
- Contract Group or LOS
- Individual Starting Point
- Prescribed Sea Tour
- Normal Shore Tour

In addition, there are "exception variables" like the following:

- Special Unit or Activity Tour Credit
- Unit Deployment/Employment Status
- Early Release Programs Shipboard Operational Holds
- Permanent Change of Station (PCS) Funding Constraints
- Sex
- Voluntary Shore Duty Curtailment

Still other considerations could be cited, but the above are sufficient to indicate some of the complexity in planning sea/shore rotations and/or evaluating policy or constraint alternatives for their sea/shore rotation consequences.

In order to supplement and/or support a manually operated (or "stubby pencil") system, several unsuccessful attempts to model and computerize the sea/shore rotation process were undertaken. The first such effort, called the Dynamic Flow Model, represented an attempt at simulation modeling by the Navy Personnel Research and Development Center (NPRDC) in the early 1970's. The model could not handle a sufficient number of the essential variables, and it was apparent that efforts to extend and enhance these capabilities could only result in an unwieldy model.

A second model called the Discounted Cash Flow (DCF) Model was developed at the Office of the Chief of Naval Operations (OP-01) in an effort to deal with "overall" issues of policy. At this level, the model failed to include sufficient detail to provide any real insight into the rotation problem.

A third effort undertaken in considerable detail by the Center for Naval Analyses (CNA) was completed in May of 1979. Called by a variety of names -- CNA Model, Expanded Sea/Shore Rotation Model, ROTATIONMOD -- this model was accepted by the Navy after a series of test runs. Partly as a result of changing personnel and partly as a result of subsequently discovered deficiencies, further work had to be undertaken in order to make this CNA model operational. B-K Dynamics, Inc., was retained for this work and, in September of 1982, completed a user's guide. This CNA model, implemented on an IBM/370 proved to be slow, expensive and confusing to use. Quoting from [5]¹, as authored by B-K Dynamics, Inc.,

"The...system is expensive.... Please keep use to a minimum, calling up the model only when a course of action is mapped beforehand and a computational strategy designed. This will cut down on its cost, which could be surprisingly high when the system is used extensively."

It was against this background of preceding research efforts that the work on GPSSR was undertaken. More than a system for effecting Sea/Shore rotation was intended. By agreement with the Navy and the CCS the model was to be able to deal with rotation scheduling in requisite detail and also lend itself to policy evaluation at more global levels. It was also to provide a basis for improved planning of "officer based" as well as "enlisted based" systems. Finally it was to provide a possible

¹ Numbers in square brackets are keyed to the references at the end of this report.

approach for integrating both officer and enlisted personnel planning to the extent that this might be feasible.

Prior experience with large and complex personnel planning models made it clear that two important types of difficulties were to be anticipated in any model that might be synthesized. First, a variety of conflicting objectives were likely to be encountered so that some way was needed for dealing with the difficulties that such conflicts can cause for most types of mathematical models. "Goal programming" was initially developed by A. Charnes and W. W. Cooper (in collaboration with R. Ferguson [2]) in order to deal with such conflicts for use on Navy personnel problems. Subsequently extended by A. Charnes and W. W. Cooper (in collaboration with R. Niehaus [3] and [6]) it also has the capability of showing where (and in what amounts) the conflicts are causing deviations from prescribed goals and policies.

The class of goal programming models thus provided an attractive basis for the combinations of rotation scheduling and policy evaluation that were wanted. This was one reason for selecting a goal programming approach to Sea/Shore rotation. Another is that it lends itself to the kinds of extensions that might subsequently be effected to "officer based" as well as "enlisted based" systems.

A second class of difficulties was also to be anticipated in the form of computational algorithms and computer codes that might be used for

these models. Ordinary goal programming computer codes would not be up to the performances required in these applications. Past experience with computer codes of "network varieties" has shown that these types of codes can now accomodate problems of huge size and complexity, provided the problems can be given characterizations that lend themselves to network representations. Again, A. Charnes and W. W. Cooper (in collaboration with R. Niehaus [3] and [6]) had previous experience and success in joining network and goal programming models in a single goal programming/network representation that could be handled by available network codes.

In the present case (as was also anticipated), still further extensions of all of these previous developments were likely to be required. The nature of these extensions are described in the sections that follow. For clarity, attention is confined to "enlisted based" Sea/Shore Rotation applications.

3.0 MODEL AND SYSTEM OVERVIEW

The model uses a goal programming network form for representing the flows of personnel within a detailing community (DC) over time. A network, being a collection of nodes and arcs, can be used to represent states and the relations of flows between them. In GPSSR, the arcs are used to represent flows of personnel between the nodes, while the nodes represent different personnel categories and status. The categories are defined by those qualifications which are needed to capture the essence of the rotation problem, and it is the number of these qualifications which directly affects the problem size. Size is not usually a problem since the software developed by the CCS is presently capable of handling several thousand nodes and tens of thousands of arcs. Some understanding of the model is required, however, since the introduction of additional parameters can affect the size of the problem in different ways, according to the strategy of representation used. For this reason, the model is set forth in Appendix A.

Currently, five qualifications are used to define the nodes: These are an individual's (1) paygrade; (2) time on tour; (3) length of service; (4) type of duty; and (5) the specified year of the planning horizon. The objective of the analytical model is to minimize the total dollar costs and goal costs, subject to certain constraints and network relations. These costs and constraints are discussed below.

Limitations and preferences for various types of personnel movement are rendered in the form of constraints and prescribed goals to reflect given rotation policies. Additional constraints include the transition rates which represent the historical rates of promotion, accession, loss, etc. of personnel. Two different kinds of goals are involved: (1) Those expressing the desire to fill billets; and (2) Those expressing the desire to rotate personnel in accordance with Navy priorities. Goal costs are assigned to reflect the relative importance of meeting these goals. Goals are derived from input to the model in the form of numbers of future personnel authorizations, or proposed changes in end strength, these changes being specified as numbers or as percentages of current staffing levels. The model also incorporates the real dollar costs associated with the Permanent Change of Stations (PCS) involved. Both real and goal costs are reflected in the minimizing objective as noted in the preceding paragraph.

The GPSSR system consists of five modules: (1) a data extraction component; (2) a transition rates module; (3) a network generator; (4) a network optimizer; and (5) a report generator. Mathematical details are supplied in Appendix A to this report. The operation of GPSSR may be summarized as follows: First, the model extracts raw data from the Enlisted Master Record (EMR). Then, after providing some automated data correction--as well as the facility for manual data adjustment--it determines smoothed historical transition rates. When these transition rates have been reviewed and respecified, a network is generated. A network code

then computes the optimum flows on this network to minimize goal costs as well as dollar costs. Finally, the system provides a report generator to display various aspects of this optimal solution in order to facilitate monitoring and/or redirection of these efforts.

4.0 GPSSR SYSTEM MODULES

4.1 EXTRACTION AND SEPARATION OF DATA

The first task is extraction of the relevant information from the Enlisted Master Record (EMR). The raw data for this purpose are currently available from the Center for Naval Analyses (CNA) in the form of magnetic tapes. Each tape contains data for several detailing communities (DCs) which need to be separated by DC for use in this model. The EMR contains a very large record for each individual from which only a few data fields are needed. After these fields have been extracted, a DC-specific file is produced containing a reduced individual record for each member in the DC.

As is true for many data sources, the EMR may (and generally does) contain some errors that need to be detected and corrected. As a result, the separation programs include an elaborate structure of error-checking to guarantee "clean" reduced DC files. The checking is accomplished, in part, through use of the many fields of overlapping information found within the EMR. Some of the checking cannot be done automatically, however, because of the many different kinds of errors potentially to be found in the tapes, and so the programs are designed to enable an operator to apply his or her own knowledge and judgment when such situations arise. Even so, this is a tedious effort, requiring some experience with the programs as well as knowledge of the nature of the data in the EMR tapes. The separation module is independent of the other modules; hence, once a

satisfactory separation and error correction have been achieved, the remaining modules may be run repeatedly for parameter studies without having to re-extract this data.

4.2 CALCULATION OF TRANSITION RATES

Having described the process for extracting and preparing the data, we now turn to the second of the five GPSSR modules, the transition rate module, which computes smoothed Markov rates for use in the constraints for the network.

4.2.1 OBTAINING TRANSITION TOTALS

The second module of the GPSSR package calculates the historical rates of accessions, losses, promotions, demotions and rotations in the years for which the data are supplied. This is done by examining the extraction from the EMR for two successive years, finding the rank and type duty for individuals in the DC both years, and thereby determining how many individuals were promoted, demoted, rotated, etc. Individuals found in only one of the two years are treated as accessions or losses to the DC.

4.2.2 CALCULATING THE SMOOTHED RATES

After determining the rank and type duty for individuals, and how personnel were transferred, this module takes the historical transition totals and calculates the (Markov) transition rates for the time span covered by the model. This is accomplished via an exponential smoothing algorithm which uses either a user-supplied smoothing factor, or, if the user prefers, a smoothing factor, α , which is stored in the computer. The exponential smoothing algorithm used is described in Appendix B. The existence of these transition rates, as reflected in the proportionality constraints, or "side constraints" of the model, would normally preclude solution by a pure network program; however, by relying on a new method of approximating these constraints, GPSSR can take advantage of the very fast pure network codes available at the CCS. The new method is explained in detail in a later section.

4.2.3 USER INTERACTION

The system is built so that the transition rates mentioned above can be modified by the user for those cases where it is known or expected that historical transition rates will not reflect the actual course of events. In particular, Enlisted Community Managers generally have access to the planned number of accessions for their community, a number which may be at variance with the historical rates--e.g. in recent years, some DCs have experienced significant expansion in size. For these, the historical rates of accession will not be a reasonable indicator of the actual accession rates observed. When this occurs, or in other like situations, the user can input the planned accessions, overriding the system-calculated rates. This interaction capability is currently being upgraded for greater "user friendliness", which will include system supplied prompts and menus to aid users in their choices.

4.3 NETWORK GENERATOR

Having described the modules for extracting the data from the EMR, and for computing the smoothed transition rates reflecting the historical proportions of promotion, demotion, etc., we now turn to the third module in the package, the network generator. This module introduces upper and lower bounds that limit the personnel flows between nodes in the network. The lower bound stipulates a minimal amount of flow that must be attained on the arc to which it applies while the upper bound provides a

capacity limit which the flow cannot exceed. The introduction of these upper and lower bounds changes the model from a pure (or uncapacited) network to one that is formally characterized as having a capacitated network structure.

Ordinary network computer codes must be modified to deal with networks that are capacited. GPSSR must also handle transition conditions that involve additional "side" constraints so that still further extensions of these network codes are required. We have avoided the use of general purpose algorithms for networks with side constraints--often called a "constrained network"--because these algorithms are not efficient for large models of this type. For models as large as ours (for a typical DC, a network with several thousand nodes and tens of thousands of arcs is generated,) use of these algorithms requires solution times which are prohibitive.

4.3.1 GENERATING BOUNDS

To achieve better solution times, GPSSR uses a new algorithm developed from our research which is designed to approximate this "constrained network" by a "pure network" which is also capacited. (A more mathematical description of these types of networks is provided in Appendix A.) This is done in two steps: First, a projection routine calculates an exact flow on each arc based on the historical transition rates generated by the previous routine, possibly modified by an informed user. This

projection provides an estimate of the flow for the entire period covered by the model. If the user is satisfied with such a quick estimate, and does not require any optimization, it is possible to proceed directly to the report generator. If, however, the user wishes to determine the flow of personnel which will "come closest to meeting goals and priorities" at minimal cost, this projection will then be embedded in a constrained network.

Part of the flexibility and efficiency of GPSSR comes from using this projection as a starting point for developing a constrained network. A user-supplied or default flexibility parameter, δ , is applied to the projected flows to generate upper and lower bounds, thus allowing flow to occur only within these bounds on the arcs to which they apply. The resulting network is a pure capacitated network. For small δ the constraints are satisfied to within a good approximation, even where the flows within the indicated bounds do not satisfy the proportionality constraints exactly. Arcs having a "window" defined by such upper and lower bounds are called "valve arcs." The flexibility parameter may be varied across the different types of arcs, so that windows of different sizes can be generated as needed.

The approach, as described to this point, confines the model to windows determined by the historical rates and flexibility parameter(s). This can be inappropriate for many planning and evaluation situations. Hence, provision is made for the addition of "bleeder" arcs to permit

deviations from historical rates, but only with a penalty cost. It is also possible to maintain rigid (historical rate) constraints, where these are known, by setting δ to 0, while putting prohibitive penalties on the "bleeders." A more detailed description is deferred to Appendix A.

While GPSSR allows a great deal of user intervention, it is designed so that it does not place heavy burdens on the user; on the contrary, very little user interaction is required. The user need only call the appropriate optimization routine and (optionally) supply flexibility parameters. The rapid solution capabilities of the optimization algorithm make it feasible to explore a variety of alternatives with different parameters. Furthermore, planned enhancements of the model's user interface will largely automate this process.

4.3.2 ATTACHING GOALS AND COSTS

By this point, a network has been generated using the historical transition rates as modified by the user's knowledge and experience. A network is thus obtained with arcs which describe every possible transition from one paygrade, length of service, and type of duty to some other possible combination of paygrade, length of service, and type of duty. On each of these arcs, we have also imposed upper and lower bounds which allow flexibility from the historical proportions. For purposes of optimization, it is then necessary to attach dollar costs and goal, or priori-

ty, costs to all these arcs, so that it makes sense for the program to optimize these costs. It is then possible to obtain the set of flows which minimizes the weighted deviations from the stated goals at the least possible dollar cost while remaining within the constraints.

At the present time a file has already been written with a set of goal and dollar costs. This has been done so potential users can experiment with the code and provide possible guidance for further directions of development. Such users will find that the file is already able to provide at least minimal automatic support for situations in which the user does not wish to supply information in the requisite detail. Users who wish to do so, however, can insert additional information about costs, goals, and priorities before running the code. Goals, in the form of manning requirements must be provided to the model at this time, and will be used to write "goal arcs." A sample file is available, so the user can see the proper format for specifying the desired billets. A description of the "goal arcs" used to represent these manning requirements is in Lovegren [4], and is further described in Appendix A.

4.4 NETWORK OPTIMIZER (VICNET)

We have described the data extraction, computation of the historical transition rates, and the generation of a capacitated network with "valve arcs," "bleeder arcs," and "goal arcs," and now provide a brief description of the network optimizer. As compared with the code described

in Lovegren [4], the current version has achieved another order of magnitude increase in speed. This is significant in its own right, and it is also indicative of the progress that continues to be made in reducing these running times. Lovegren's work reduced the running time for solving the sea/shore rotation problem from 24 hours to one hour; the current version runs in about 40 seconds for a DC. Furthermore, the previous version did not take into account "real dollar" costs, as does the current version. To distinguish between the real and goal costs, the new code uses an approach ¹ that first minimizes deviations from stated goals, then achieves this result at the lowest possible dollar cost. In addition the model is now capable of keeping track of different kinds of dollars, which can be important when funds are earmarked and non-transferable.

4.5 REPORT GENERATOR

This section presents a summary of the reports which may be obtained from the GPSSR system. These output modules were developed concurrently with the modules for extracting data, so, while data have been extracted from the Navy's EMR for actual detailing communities, these report modules were tested on hypothetical data, and the charts and tables presented here

¹ This corresponds to what is technically called a "non-Archimedean" approach as described in detail in [1], and as is briefly described in Appendix A.

are intended only to show the sorts of reports which may be generated, not the results from a real detailing community.

Chart 1 was obtained by downloading data from the network optimizer to an IBM PC, then graphing this data with the LOTUS 1-2-3 program. All 6 duty types are presented, i. e. CONUS shore duty, arduous sea duty, over-sea duty, non-rotated sea duty, neutral duty, and oversea preferred duty. A copy of the instructions, or template, for the IBM PC is available with the GPSSR system, although the user must provide a copy of the LOTUS 1-2-3 program in order to use this template.

Chart 2 accumulates all sea and shore duty so the user may see the overall Sea/Shore picture. This chart is automatically generated by our template for the LOTUS 1-2-3 program from the same data which produced Chart 1.

Billets vs. Inventory

Sea/Shore Rotation




Chart 1.

Billets vs. Inventory

Sea/Shore Rotation




Chart 2.

The data presented by Charts 1 and 2 may also be obtained in tabular form. However, the data used for the accompanying table is not the same set of hypothetical data that was used for Charts 1 and 2, since the programs to generate the tables and charts were being developed in parallel. In a production environment, the data from Table 1 would be downloaded from a mainframe computer to an IBM PC, or some other personal computer, and input to our template for the LOTUS 1-2-3 program, or to some other program with similar capabilities to produce Charts 1 and 2.

The tabular form presents, in addition to the information in the charts, details about any combination of scheduled (i.e. expected under an optimization program) promotions, demotions, accessions, losses and rotations for all the years covered by the model. Table 1 presents a sample of these capabilities. From the Table, we have extracted the page presenting CONUS shore duty, arduous sea duty, oversea shore duty, and non-rotated sea duty for the final period of a sample run. The information on the inventory scheduled by the optimizer, the user's goals, and the deviations from those goals is always presented. In addition, the user requested information on promotions, losses, and accessions. The total movement of personnel also includes demotions and rotations, which were not requested for this run of the report generator but which can also be displayed. As a result of this flexibility, the vertical columns do not sum to the total inventory unless all categories of movement are displayed.

PERIOD: FY46

PAYGRADE

STATED ROTATION POLICY

TYPE	DUTY:CONUS SHORE DUTY				E9	E9	TOTAL
	E4	E5	E6	E7			
INITIAL INVENTORY	2	2	2	2	3	3	2289
PROMOTIONS INTO PAYGRADE	131	404	518	234	117	81	137
LOSSES OUT OF PAYGRADE	0	134	151	0	0	0	24
ACCESSIONS INTO PAYGRADE	5	13	0	0	0	0	186
SCHEDULED INVENTORY	30	60	30	17	9	4	1315
PAYGRADE GOALS	275	1130	609	306	106	85	1614
DEVIATIONS FROM GOALS	72	539	444	251	113	80	1631
	1252	591	165	55	-7	5	

PERIOD: FY46

PAYGRADE

STATED ROTATION POLICY

TYPE	DUTY:AROUNDUS SEA DUTY				E9	E9	TOTAL
	E4	E5	E6	E7			
INITIAL INVENTORY	31/3	2757	1486	984	422	97	8993
PROMOTIONS INTO PAYGRADE	0	8	21	37	0	0	66
LOSSES OUT OF PAYGRADE	2	6	10	1	0	0	19
ACCESSIONS INTO PAYGRADE	1274	653	159	52	35	4	2146
SCHEDULED INVENTORY	4345	2956	1120	1158	451	77	10612
PAYGRADE GOALS	2639	2224	1641	501	379	84	7378
DEVIATIONS FROM GOALS	1706	732	-121	597	72	41	3056

PERIOD: FY46

PAYGRADE

STATED ROTATION POLICY

TYPE	DUTY:OVERSEA SHORE DUTY				E9	E9	TOTAL
	E4	E5	E6	E7			
INITIAL INVENTORY	1	6	9	10	6	4	37
PROMOTIONS INTO PAYGRADE	0	0	0	0	0	0	0
LOSSES OUT OF PAYGRADE	0	0	0	0	0	0	0
ACCESSIONS INTO PAYGRADE	1	0	1	0	0	0	2
SCHEDULED INVENTORY	3	5	9	11	4	4	39
PAYGRADE GOALS	1	2	11	6	2	3	32
DEVIATIONS FROM GOALS	2	3	-2	5	-1	1	7

PERIOD: FY46

PAYGRADE

STATED ROTATION POLICY

TYPE	DUTY:NON-RATED SEA DUTY				E9	E9	TOTAL
	E4	E5	E6	E7			
INITIAL INVENTORY	467	361	213	126	61	8	1249
PROMOTIONS INTO PAYGRADE	0	0	0	0	0	0	0
LOSSES OUT OF PAYGRADE	0	0	0	0	0	0	0
ACCESSIONS INTO PAYGRADE	197	90	17	7	5	0	316
SCHEDULED INVENTORY	542	388	212	131	62	9	1358
PAYGRADE GOALS	396	299	195	81	45	5	1012
DEVIATIONS FROM GOALS	146	87	17	50	17	4	326

Table 1.

Table 2 contains a collage of the larger printout from which the report shown in the previous table was extracted. This printout shows the optimal rotation policy broken down by year, type of rotation, paygrade, time on tour, and length of service. Several sections have been pasted together to present a better view than was possible from any one section. The arc numbers and names indicate the different sections from which they were extracted, as explained below.

At the top of the table is shown, as the problem title, the name of the community covered, the flexibility option, delta, and the smoothing factor alpha. The value $\text{delta}=0$ shown here means that the user did not use the flexibility option, and the smoothing factor $\text{alpha} = 0.2$ was used to project the transitions covered by the exponential smoothing formula of Appendix B in this case. Finally, the total goal deviation resulted in penalties of 15560 from the goal deviation penalties used, and \$280,020 is the estimated (best) PCS cost associated with the program for which the details in Table 2 form a part.

Reading from left to right the column headings refer to the following: ARCNUMBER

This is the number of the arc as it was read into the network optimizer. We have presented a selection of the first few arcs, and three other sections taken from the 2000s and 8000s. The first few arcs represent initial supply, the 2000 arcs represent rotation arcs (with positive dollar cost), the first set of 8000 arcs represent goal arcs, with positive penalty costs, and the second set of 8000 arcs represent the arcs which connect the goal arcs back to the beginning of the network to form a complete circuit.

FROM NODE

These are the source nodes from which each arc originates. The code tells the type of arc, paygrade, length of service, etc.

TO NODE

This is the destination of the arc. 1P3 02 means (in order) year 1 of the optimization, promotion arc (P), paygrade E3, length of service less than 1 year (0), and type duty 2 (arduous sea.)

GOAL COST

Penalty assigned per unit flow on this arc. We have put goal costs in this formulation only on failure to meet desired personnel levels, with the -5 indicating that a cost of 5 units is assigned to falling below the requirement, and the 1 indicating a cost of 1 unit is assigned to exceeding the desired level. These costs are not expected necessarily to reflect the desires of actual users. Also, upper and lower bounds of 0 on the arc with -5 unit cost indicate

that 0 personnel were desired for this category. This is because, for this run, no goals were assigned, so the program used 0 for all the goals. The small cost of exceeding the goals, coupled with the lack of flexibility, caused personnel to be scheduled into the usual categories anyway.

DOLLAR COST

PCS cost per person assigned.

UPPER BOUND

Maximum flow allowed on the arc. Since no flexibility was allowed ($\Delta = 0$) this will be equal to the lower bound, forcing the flow to be equal to the set upper (or lower) bound, except on certain "goal" arcs, where violations are penalized but not prevented.

LOWER BOUND

The minimum flow allowed on the arc.

ARC FLOW

Actual flow on the arc.

ARC COST (G)

Flow multiplied by goal cost.

ARC COST (\$)

Flow multiplied by dollar cost.

MARG COST (G)

Marginal cost, i.e. the penalty incurred by sending one more person along arc.

MARG COST (\$)

Dollar cost incurred by one more person along arc.

PROBLEM TITLE COMMUNITY: 4000 DELTA= 0.00 ALPHA= 0.20

OPTIMAL SOLUTION

TOTAL NODES= 9916 ARCS= 8514

GOAL OBJ FCN= 15560 DOLLAR OBJ FCN= 280020 NO. ITER= 2932

ARC NUMBER	FROM NODE	TO NODE	GOAL COST	DOLLAR COST	UPPER BOUND	LOWER BOUND	ARC FLOW	ARC COST(1)	ARC COST(2)	MARG COST(1)	MARG COST(2)
1	0A0000	1P3 02	0	0	71	71	71	0	0	0	0
2	0A0000	1P3 04	0	0	1	1	1	0	0	0	0
3	0A0000	1P3 11	0	0	13	13	13	0	0	0	0
4	0A0000	1P3 12	0	0	828	828	828	0	0	0	0
5	0A0000	1P3 14	0	0	168	168	168	0	0	0	0
6	0A0000	1P3 21	0	0	10	10	10	0	0	0	0
2826	1W7114	257124	0	10	10	10	10	100	0	0	10
2827	1W7115	257125	0	10	5	5	5	50	0	0	10
2828	1W7121	257132	0	10	7	7	7	70	0	0	10
2829	1W7121	257131	0	10	19	19	19	190	0	0	10
2830	1W7122	257131	0	10	16	16	16	160	0	0	10
2831	1W7122	257132	0	10	44	44	44	440	0	0	10
11	0A0000	1P3 34	0	0	32	32	32	0	0	0	0
12	0A0000	1P3 41	0	0	6	6	6	0	0	0	0
9338	0G3001	0G3001	-5	0	0	0	0	0	0	-6	0
9349	0G3001	0G3001	1	0	250000	0	304	304	0	0	0
9390	0G3002	0G3002	-5	0	0	0	0	0	0	-6	0
9391	0G3002	0G3002	1	0	250000	0	4365	4365	0	0	0
9392	0G3003	0G3003	-5	0	0	0	0	0	0	-6	0
9393	0G3003	0G3003	1	0	250000	0	3	3	0	0	0
9394	0G3004	0G3004	-5	0	0	0	0	0	0	-6	0
9395	0G3004	0G3004	1	0	250000	0	542	542	0	0	0
28	0A0000	1P3 64	0	0	3	3	3	0	0	0	0
29	0A0000	1P3 66	0	0	4	4	4	0	0	0	0
75	0A0000	1P3 71	0	0	7	7	7	0	0	0	0
875	0G7001	0Z0000	0	0	25000	0	306	0	0	0	0
9477	0G7002	0Z0000	0	0	25000	0	451	0	0	0	0
9478	0G7003	0Z0000	0	0	25000	0	6	0	0	0	0
9499	0G7004	0Z0000	0	0	25000	0	62	0	0	0	0
9500	0G7005	0Z0000	0	0	25000	0	37	0	0	0	0
9501	0G7006	0Z0000	0	0	25000	0	10	0	0	0	0
9502	0G8001	0Z0000	0	0	25000	0	106	0	0	0	0
9503	0G3002	0Z0000	0	0	25000	0	125	0	0	0	0
9504	0G3003	0Z0000	0	0	25000	0	1	0	0	0	0
9505	0G3004	0Z0000	0	0	25000	0	14	0	0	0	0
9506	0G4005	0Z0000	0	0	25000	0	8	0	0	0	0
9507	0G5005	0Z0000	0	0	25000	0	6	0	0	0	0

TABLE 2.

5.0 MONITORING FEATURES

5.1 DISPLAYING STATISTICS FROM THE EMR


Once the data have been extracted from the EMR, the University of Texas computer system provides an advanced graphics facility which makes it possible to monitor past and present activities and consequences of personnel management as reflected in the EMR data. This is an important function of the system, because the size and complexity of personnel transfers, as well as the existence of numerous exceptions often masks the real situation from managers trying to obtain a good picture with only manual methods for information extraction and summarization from the data. For example, important topics like how much of an existing "rotation policy" is actually being implemented in view of the exceptions need to be addressed regularly.

As a start toward developing desirable monitoring capabilities, GPSSR currently employs the statistical package (SAS) to calculate and display various statistics concerning the data. The package can produce a graph of almost any combination of the variables found in the EMR. GPSSR provides the ability, using a single command, to generate those charts and graphs deemed useful for policy analysis. The following examples demonstrate a few aspects of this ability.

The first set of graphs, represented in Figures 1 and 2, is obtained before any modelling or optimization has been done. In principle, the user could obtain these graphs by extracting from the data the record of every individual in a DC, then making these graphs manually by plotting such things as time on tour vs. type duty with a "stubby pencil" on graph paper. As part of our GPSSR modeling project, however, we have completely automated the process, so that, with a single command, the user can see these results for purposes of monitoring the status of the current implementation of Sea/Shore rotation policy and to better plan future rotation strategies. For example, and just as an example, we have chosen to display a bar chart, showing the distribution of time on tour for two of the six types of duty in DC 4000 at Length Of Service 5-17 in the 1982-1983 time frame. The two types of duty shown are (1) CONUS shore duty, in Figure 1; and (2) Arduous sea duty, in Figure 2. For each duty type, we display a histogram showing the percentage of the community who have spent 1, 2, 3, 4, or 5+ years assigned to that type duty so that one can observe how far along their tour most of the community lies, and where, consequently, an extension or shortening of tour length will have the most effect.

TYPE DUTY=1

PERCENTAGE




LENGTH OF SERVICE BETWEEN 5 AND 17

Figure 1.

TYPE DUTY#2

PERCENTAGE




LENGTH OF SERVICE BETWEEN 5 AND 17

Figure 2.

As another example to indicate some of the possibilities of such graphic capabilities, we provide the following plots of time on tour vs. paygrade in Figures 3 and 4. The dotted lines show the distribution of 90% of the community, while the solid line shows the mean over all lengths of service. Where the two dotted lines diverge very markedly, the averages are not sufficiently meaningful for drawing firm conclusions. Conversely, when the dotted lines lie close to the mean, little information is lost by using an average as opposed to considering all the observations separately. Again, we show these plots for type duties 1) and 2), CONUS shore duty and arduous sea duty, respectively.


TYPE DUTY=1



LENGTH OF SERVICE BETWEEN 5 AND 17

Figure 3.

TYPE DUTY=2



5.2 ADDITIONAL MONITORING CAPABILITIES

Not immediately available from the EMR are rates of accession and loss, and length of tour as opposed to time on tour. In order to obtain these rates, we had to compare two years of the EMR. Note that accession and loss, for our purposes, refer to a single community. People who transfer from one DC to another are considered an accession to their new community and a loss to their old community. For purposes of filling a given community's billets, this should not be an unreasonable definition.

These data, as well as promotion and demotion rates are available in a readable file, and plans exist for a report generator that will make them even more accessible. In addition, a graph package is planned that will present the data in a form similar to the example shown. The example was prepared using the SAS package. However, as part of the continuing effort to develop an intelligent user interface, a more user-friendly plot interface is planned which will be much easier to access than the SAS plot package.

As a further example of the kinds of GPSSR graphs that can be generated at this point, the chart in Figure 5, based on hypothetical data, is presented to show the proportion of personnel promoted while on type duty 2 as a function of paygrade. Here the dotted lines show the range of 90% of the data, averaged over length of service and time on tour.

PERCENTAGE PROMOTED VS PAYGRADE FOR TYPE DUTY 2 DETAILING COMMUNITY 4000




Figure 5.

6.0 WORK IN PROGRESS

Parallel efforts are also currently under way in the CCS to improve the GPSSR capability and performance. A brief description of some of these follows : -

6.1 DEVELOPMENTS IN THE THEORY

- Introducing an alternative goal concept, that of goal "length of tours." To that end, the time on tour has been introduced as another node dimension, enabling the model to calculate penalties based on deviations of desired lengths of tour. These penalties are then added in with the other goal costs, representing deviations from the planned billets, which were already in the model. Early rotations, which might be of concern to the DC personnel management, are penalized. Likewise, late rotations, which might cause individuals to quit, are also penalized.
- Studying the effects of the non-Archimedean optimization on the rate of change in the DC strength. As explained in 3.4, the minimization is taken first on the goal costs and only then over the real costs. If the end strength goals are somewhat higher than the start inventory, the model, given only the dollar costs for maintaining personnel, and only goals for strength in the final year, will try to access people as late as possible to avoid the costs of carrying them along the network. The computed solution may then suggest abrupt changes in manning for the DC, all taking place in the last year under consideration. However, the introduction of intermediate goals via "valve arcs" and "bleeder" arcs, will cause the model to provide for gradual changes and smooth-out possible saw-like jumps in the personnel curve.
- Considering different scenarios and objectives regarding the male and female personnel in certain DCs. Special attention needs to be devoted in the modeling process to address problems resulting from legal constraints and the lack of available positions at sea which are adequate for women (e.g. older ships must be modified to accommodate female personnel). This situation creates imbalances in the rotation policies applicable to different sexes.

6.2 TECHNICAL ENHANCEMENTS/REFINEMENTS

- Developing an "intelligent" user interface for the GPSSR. One of the main goals of this effort is to provide powerful interactive capabilities, so that a decision maker need be neither a computer expert nor an operations research expert in order to use the system. Using normal Navy language, the user should be able to explain his problem to the system, which will automatically call the appropriate programs, prompt the user for specification of parameters and directives, and produce the desired output. Natural language processing in all detail is more ambitious than we expect to achieve, but we do intend to push very far in that direction.
- Producing more summary reports as derived from the global output file. One such report should aggregate the costs resulting from the personnel movement in the network. Currently, the aggregation is by paygrade and type duty, calculated separately for the different types of costs.
- Enhancing the quality of the input data. We hope to improve our understanding of how to handle some of the inputs which have not, as yet, been thoroughly checked. The quality of cost information, for example, must be improved.

7.0 SUMMARY AND CONCLUSIONS

7.1 SUMMARY

The GPSSR system is a sophisticated Management Information/ Decision Support System, produced by the Center for Cybernetic Studies at the University of Texas for the U. S. Navy. The system handles possibly contradictory information by optimizing, via goal-programming, over suitable goals, using a capacitated network model structure with computation orders of magnitude faster than that of previous Sea/Shore rotation models. The system contains a monitoring capability which provides a manager with previously unavailable information about the Sea/Shore rotation policy actually implemented.

7.2 CONCLUSIONS AND DIRECTIONS OF FURTHER WORK

This system will be a useful tool for Navy managers and planners. It is also general enough to be applied in solving an array of problems other than the sea/shore rotation problem. It can be used to solve any problem--including optimization problems--with elements involving goals and flows, capacities and costs. Its goal programming features permit identification and analyses of deviations from goals caused by one or more of these elements such as might be involved in officer or enlisted based problems and the planning of optimal force structures. Finally, the dual evaluators are available for exploitation in policy analyses and evaluations such as are likely to be present in allocation policy problems associated with manpower planning.

APPENDIX A

MATHEMATICAL DESCRIPTION OF THE MODEL

Much of the following is abridged from Lovegren (for a fuller explanation see [4]). A network may be visualized as a collection of nodes $S = \{1, 2, \dots, n\}$, and between these nodes a set of arcs. Along each arc is a flow x_{ij} , the flow from node i to node j . If $x_{ij} < 0$, this represents a flow of $|x_{ij}|$ from node j to node i . Using c_{ij} to represent the cost per unit flow from node i to node j , the pure network optimization problem is then

$$\min_{x_{ij}} \sum_i \sum_j c_{ij}(x_{ij}) \quad (A-1)$$

subject to the network constraints "lo que entra sale" or "what comes in goes out," i.e.,

$$\sum_i x_{ik} - \sum_j x_{kj} = a_k \quad (A-2)$$

This says that, at each node k , the total flow going into the node minus the total flow going out of the node is equal to the net inflow or outflow at that node. In matrix notation, (A-1), (A-2) can be written

$$\begin{aligned} \min \quad & \underline{c}^T \underline{x} \\ \text{s.t.} \quad & N\underline{x} = \underline{a} \end{aligned} \quad (A-3)$$

Components of the \underline{c} vector represent the cost per unit flow on each arc and the component of the \underline{x} vector represent these flows (from node i to node j). Since every arc must go between two nodes--into one and out of the other--each column of the N matrix has precisely two non-zero entries: $+1$ and -1 . All other entries in each column are zero except for these ± 1 values which are incident on nodes i and j , respectively. The N matrix is called a node-incidence matrix. In fact, any matrix with this property may be considered a node incidence matrix, with

each row associated with a node for which non-zero entries appear. Each column represents an arc with the ± 1 values indicating the nodes on which it is incident. That is, since the column has a +1 in row j, say, and a -1 in row i, it may be graphically represented as an arc from node i to node j.

Additional constraints of the following type may be added to form a (pure) capacitated network problem:

$$l_{ij} \leq x_{ij} \leq u_{ij} \quad (A-4)$$

The complete problem is then

$$\begin{aligned} \min \quad & \underline{c}^T \underline{x} \\ \text{s.t.} \quad & N\underline{x} = \underline{a} \\ & \underline{l} \leq \underline{x} \leq \underline{u} . \end{aligned} \quad (A-5)$$

Note that (A-5) looks similar to the general linear programming (LP) problem

$$\begin{aligned} \min \quad & \underline{c}^T \underline{x} \\ \text{s.t.} \quad & A\underline{x} = \underline{b} \\ & \underline{x} \geq 0 \end{aligned} \quad (A-6)$$

where A is an arbitrary matrix. However, the algorithms to solve (A-6) require some two orders of magnitude more computations than the algorithms to solve (A-5). In particular, the Center for Cybernetic Studies has developed one of the most efficient network optimizers available, a specialized package which can solve (A-5) but not (A-6), and is two orders of magnitude faster

than a general purpose LP package such as MPSX when applied to an optimization of the form (A-5).

REPRESENTATION OF NETWORKS

Given a problem in the form (A-5) it is easiest to visualize the problem by drawing it as a network. An arc of the network is represented as in Figure (A-1), where c_{ij} is the unit cost associated with the arc, l_{ij} is the lower bound, or minimum flow required on the arc, and u_{ij} is the upper bound, or maximum flow which may be allowed. (It is impossible to display the entire GPSSR network, as it contains over 10,000 nodes and 20,000 arcs. Enough simplified subsections are presented below to give a good picture of the entire network.)




FIGURE A-1

GOAL PROGRAMMING

The GPSSR program tries to meet personnel goals as closely as possible. This implies that the objective function contains terms of the form

$$\min_{x_{ij}} \sum_i \sum_j c_{ij} |x_{ij} - g_{ij}| \quad (A-7)$$

where g_{ij} is the goal and the vertical strokes represent an absolute value of the difference between g_{ij} and x_{ij} .

Minimizing this function gets the x_{ij} "as close as possible" to the g_{ij} . The objective function (A-7), however, does not appear to be in the form required by (A-5). A transformation due to Charnes, Cooper, et al. [1] brings us back to (A-5):

$$\min \sum_i \sum_j c_{ij} (\delta_{ij}^+ + \delta_{ij}^-) \quad (\text{A-8})$$

$$\begin{aligned} \text{s.t.} \quad & x_{ij} - \delta_{ij}^+ + \delta_{ij}^- = g_{ij}, \\ & \delta_{ij}^+, \delta_{ij}^- \geq 0 \end{aligned}$$

Note that the δ_{ij}^\pm , represent deviations above and below the goal g_{ij} resulting from the flow value assigned to x_{ij} . These deviations are accorded penalties or "costs" c_{ij} per unit in the functional being minimized.

The network representation of (A-8) in terms of goal arcs is shown in Figure (A-2), and an alternative but equivalent representation in Figure (A-3).




FIGURE A-2




FIGURE A-3

NETWORKS WITH SIDE CONSTRAINTS

Let N be a node incidence matrix, P an arbitrary matrix, and consider the problem

$$\begin{aligned} \min_{\underline{x}} \quad & \underline{c}^T \underline{x} \\ \text{s.t.} \quad & N\underline{x} = \underline{a} \\ & P\underline{x} = \underline{b} \\ & \underline{l} \leq \underline{x} \leq \underline{u} \end{aligned} \tag{A-9}$$

This is a network with side constraints, the side constraints being $P\underline{x} = \underline{b}$, where P is a matrix of coefficients and \underline{b} a vector of additional conditions. In addition the \underline{l} and \underline{u} are vectors that impose lower and upper bounds on the possible choices of \underline{x} .

Algorithms to solve problems of the form (A-9) exactly are two orders of magnitude slower than algorithms to solve (A-5). For GPSSR, we assume that the rates of promotion, loss, etc. will be similar to the historical rates, and these are our side constraints. Thus, the network constraint on the flow through a node j is simply the previously discussed condition that flow in equals flow out. The proportionality constraints associated with P put an additional requirement on the flows out of the node--e.g. they must be proportional to the total flows through the node.

It is reasonable to solve (A-9) approximately. In fact, since we do not expect historical rates to be followed exactly, a more realistic version of (A-9) is

$$\begin{aligned} \min \quad & \underline{c}^T \underline{x} \\ \text{s.t.} \quad & N\underline{x} = \underline{a} \\ & -\underline{\delta} \leq P\underline{x} - \underline{b} \leq \underline{\delta} \end{aligned} \tag{A-10}$$

where the components of $\underline{\delta}$ represent the maximum and the components of $-\underline{\delta}$ represent the minimum admissible deviations from the corresponding components of \underline{b} .

By assigning a proportionality constraint to every arc, P becomes invertible. Then we may write

$$\begin{aligned} \min_{\underline{x}} \quad & \underline{c}^T \underline{x} \\ \text{s.t.} \quad & N\underline{x} = \underline{a} \\ & P^{-1}(\underline{b} - \underline{\delta}) \leq \underline{x} \leq P^{-1}(\underline{b} + \underline{\delta}) \end{aligned} \tag{A-11}$$

since P^{-1} is non-negative.

Letting $\underline{l}, \underline{u} = P^{-1}(\underline{b} \pm \delta)$, (A-11) is seen to be like (A-5) and we may employ the power of our specialized network optimizer.

NON-ARCHIMEDEAN NETWORK OPTIMIZER

The cost vector \underline{c} in (A-11) represents a set of penalties for failing to achieve goals. They are artifacts of the model and not actual dollar costs. In some scenarios, it is imperative that goals be met as closely as possible regardless of cost, but if alternate solutions exist which meet the goals equally well, then the solution which minimized real dollar costs should be chosen. This is achieved by using a non-Archimedean (or non-standard) version of the network optimizer, which solves

$$\begin{aligned} \min \quad & \underline{c}_1^T \underline{x} + \epsilon \underline{c}_2^T \underline{x} & (A-12) \\ \text{s.t.} \quad & N\underline{x} = \underline{a} \\ & \underline{l} \leq \underline{x} \leq \underline{u} \end{aligned}$$

where ϵ is a non-standard infinitesimal. In this formulation, $\underline{c}_1^T \underline{x}$ is pre-emptively" minimized which means that $\underline{c}_2^T \underline{x}$ is considered only in a way that will not alter optimal values of $\underline{c}_1^T \underline{x}$. With this formulation the network optimizer will achieve stated goals as closely as possible and then minimize dollar costs.¹

¹ Alternatively, if these are two sets of goals, one pre-emptively important, then two sets of penalty costs could be used.

THE GPSSR NETWORK

We shall shortly present a simplified numerical illustration but first we complete our network interpretation and development and introduce some additional terminology as follows:

In GPSSR, each time period's network is divided into four parts. These four parts do not represent, for example, the four seasons of a year or the four weeks of a month, but are just logical divisions of the network. The first section consists of the promotion spray arcs. These arcs handle all the promotions, demotions, losses, and accessions. Type of duty is held fixed at this point. The second section is the promotion hose arcs. The flow through a node is not computed by the network optimizer, only flow through an arc. Every node, representing a category of personnel, is then connected by a hose arc to a node representing the same category, so the flow through the hose arc allows us to observe the flow through the node, which is equal to the total number of personnel in that category during the time period in question. The third section consists of the rotation spray arcs which handle all the changes in the type of duty. Finally is a section of rotation hose arcs. A diagram is given in Figure (A-4).

The hose arcs have not been altered from the description in Lovegren [4]. However, the promotion spray arc (PR) section and the rotation spray arc (RS) sections of that network have been modified from the characterizations used by Lovegren. It is impossible to show an entire PR section or RS section for the actual network. However, we will present a simplified subsection which includes all the essential features.

We first present a simplified PR section in Figure (A-5). This section is replicated for every combination of type duty, time on tour, and length of service.






FIGURE A-5

This is still not a complete diagram of a PR subsection, as the accession and loss arcs have been omitted for purposes of simplifying the representation. The section illustrates how valve arcs and bleeder arcs are used to account for historical rates of promotion, demotion, accession and loss. Flowing into this section are the RS hose arcs, and the flow continues with the PR hose arcs, neither of which are shown. What is shown are the valve and bleeder arcs for promotion and demotion (accession and loss being handled similarly). The valve arcs allow personnel to move through the network at historical rates ($\pm\delta$) with no penalty costs. If historical rates are to be violated by more than a specified percentage, however, the bleeder arcs assign appropriate penalties which will be incurred in the violation. These penalty rates are used to discourage violations and, indeed, increasing deviations can be penalized at increasing values--although this is not illustrated in the diagram.

In the actual diagram for a hypothetical community, 45 people would have been promoted from E3 to E4 based on historical rates. Thus, the segment drawn allows 40-50 persons to be promoted with no penalty, but imposes a cost of five units per person for deviations above or below this range.

We next show a simplified version of the RS spray arcs for the last period of the problem, along with the goal arcs in Figure (A-6). We have shown only a single paygrade and length of service. Also, coming into the section are the promotion hose arcs, which are not shown. We illustrate a situation with a maximum tour length of three periods, and only three types of duty. Even so, this subsection has 18 nodes for the RS part, and 27 RS spray arcs, only 9 of which have been drawn.

We assume that policy is to keep personnel in duty type three for three years, then transfer them, if possible, to duty type two. If this is not possible, the second choice is a transfer to duty type one, with the last




FIGURE A-6

choice being for personnel to remain an additional period in duty type three. Penalties are therefore assigned to premature rotations, to late rotations, and to rotations to duty type one rather than duty type two, so that three years of duty type three would give the minimum amount of penalty charges. The relative weights are purely hypothetical, chosen for illustration only.

We have assumed that time on tour is not relevant to meeting the staffing levels indicated by the goal arcs in the diagram, so converge arcs sum over all times on tour to the converge nodes which are then connected by the terminal nodes by the goal arcs. Goal costs and upper and lower bounds exist on all the goal arcs, but are only written on the first set of goal arcs in the diagram.

A Numerical Example

In this section, we shall present, a simple illustrative example. For this, let us assume that the community is stable, i.e. that the number of personnel lost equals (approximately) the number gained through promotions, demotions, attritions, etc. Let us then concentrate on the rotation policy goals vs. staffing goals for a single paygrade i.e, we shall fill in the data for the arcs for Type Dutys 1 and 2 in Figure A-6, and solve the resulting network. Recall from page A10, that the first choice for Duty Type 3 was to rotate to Duty Type 2, and the second choice was to rotate to Duty Type 1. For our example, suppose that the first choice for Duty Type 2 is to rotate to Duty Type 1, with second choice being to rotate to Duty Type 3, and that the first choice for Duty Type 1 is to rotate to Duty Type 3, with Duty Type 2 the second

choice. These choice preferences are summarized in Table A-1.

Type duty To	1	2	3
From			
1	3rd	2nd	1st
2	1st	3rd	2nd
3	2nd	1st	3rd

Table A-1 Goals for type of rotation

The above table gives the preferred rotation sequence. The next consideration is tour length. We assume that policy is to have personnel serve 1 full period of Type 1 Duty, and two full periods of Type 2 and 3 Duty before transfer, as summarized in Table A-2.

Type duty	1	2	3
Tour Length	1	2	2

Table A-2 Goals for tour lengths

Finally, we assume that the average PCS cost is \$5000 from Duty Type 1 to Duty Type 2, \$9000 from Duty Type 2 to Duty Type 3, and \$20,000 from Duty Type 1 to Duty Type 3, as summarized in Table A-3.

Type duty To	1	2	3
From			
1	\$0	\$5000	\$20,000
2	\$5000	\$0	\$9000
3	\$20,000	\$9000	\$0

Table A-3 Average PCS Costs

These data are hypothetical, of course, and greatly simplified for ease of understanding, to only three types of duty and three periods for tour length. The part of these data for Type Duty 3 are drawn in the network fragment of Figure A-6, on page A11. An extract of a single arc of this figure is presented as Figure A-7. On this arc from Type Duty 3, time on tour 2 to Type Duty 2, time on tour 1, we have indicated the 2 units of goal cost and also the \$9000 PCS cost as represented in the above tables. See the circled numbers on this arc. The parenthesized values, $(0, \infty)$ mean that all flows are in the direction indicated by the arrow and there is no upper bound imposed on these flows.




FIGURE A-7

Legend: The dots indicate the presence of other node-arc incidences and penalties which are shown in detail in Figure A-6.

Again we emphasize that we are trying to restrict this discussion to simple versions of a complex problem--while recalling that our model and algorithm with associated software can handle very large problems with extremely fast solution times. Since drawing the entire network of 24 nodes and 42 arcs would only complicate our discussion for this example, the arcs for Types Duty 1 and 2 are merely indicated, rather than drawn in full.

Referring to Figure A-6, the set of arcs on the extreme right are GOAL arcs; the arcs adjacent to the GOAL arcs are the CONVERGE arcs; and the leftmost arcs are ROTATION arcs. Goal penalty "costs" must be specified for all the ROTATION arcs and GOAL arcs. We have chosen to assign a penalty of 2 units for rotation at the right time to the wrong type duty; a penalty of 2 units for early rotation to the right type duty; and a penalty of 4 units for late rotation or early rotation to the wrong type duty. We have also assigned a penalty of 2 units per person for under- or over-staffing type duty 1, and a penalty of 3 for under- or overstaffing type duty 2 or 3. [These penalties, we may note, need only be provisional. They can be used to obtain a trial solution from which we can decide whether or not to change these penalties to better reflect priorities or policy preferences, as we shall illustrate with an additional example.] A summary of the actual numbers appears in Table A-4.

Type duty	To	Time on Tour 1			Time on Tour 2			Time on Tour 3		
		1	2	3	1	2	3	1	2	3
From										
1		0	4	2	4	2	0	4	2	0
2		2	0	4	2	0	4	0	4	2
3		4	2	0	4	2	0	2	0	4

Table A-4 Goal Costs for Rotation Arcs

The above Table A-4 gives the simplified data for this hypothetical example. The penalty for "rotating" from Type duty 1 to Type duty 1, or Type duty 2 to Type duty 2 etc. (i.e., not rotating at all) during Time on Tour 1 is thus 0, since no one should be rotated with Time on tour 1. The penalty for rotating from Type duty 1 to type duty 3 during Time on tour 1 is 2 units, since it is desired that personnel on Type duty 1 rotate to Type duty 3, although not before completing one full period on tour. The penalty for rotating from type duty 1 to type duty 2 before completing one full period is 4 units, since this violates two aspects of rotation policy, rotating too soon, and to the wrong type duty. The other penalties are similarly explained, and all the penalties for this example are summarized in the Table.

Having set the costs on all the arcs in this example, we need to initialize the network with starting inventories of personnel, which we assume are as in Tables A-5. Thus, for instance, as noted in this Table, we have 10 persons with Type Duty 1 and time on tour 1; 40

persons with Type Duty 2 and time on tour 1; etc.

Type duty	1	2	3
Time on Tour			
1	10	40	15
2	10	40	15
3	10	40	15

Table A-5 Starting Inventories

While we have spread the personnel out evenly over the various tour lengths to simplify computation, in the real GPSSR the distribution of personnel by tour length will be determined by the actual data.

Finally, in order to complete the network, we need to state our desired staffing levels. As summarized in Table A-6, we have chosen as goals 35 persons in Type Duty 1, 125 persons in Type Duty 2, and 50 persons in Type Duty 3.. These numbers were chosen to reflect possible situations where goals might exceed available inventories and where it is not possible to conform to stated rotation policies in all detail. This will help to show how GPSSR could be used to resolve such conflicts.

Type duty	1	2	3
Staffing Goal	35	125	50

Table A-6 Goals for number of personnel in each type duty

Since our hypothetical inventory is less than our desired staffing levels, it will not be possible to meet all goals. The model solution obtained from GPSSR is "as close as possible," however, with the resulting divergence shown in Table A-7.

Type duty	1	2	3
Staffing Goal	35	125	50
Best Possible Schedule	<u>35</u>	<u>110</u>	<u>50</u>
Divergence	0	15	0

Table A-7 Goals vs Best Possible

Using the dollar PCS (personnel transfer) cost from Table A-3 we find that the solution displayed in Table A-7 will involve a dollar cost of \$570,000. In addition, 30 people (about 15%) had to be rotated to the second choice in the preferred rotation sequence. No one is rotated either early or late as a result of this optimization. Some of the inventory of Type 1 duty were already overdue for rotation, but all these were rotated.

A complete list of all rotations and costs is given in Table A-9, but before considering this solution, let us first consider whether the "goal penalty costs" that yielded this solution are the appropriate ones to meet imperatives in the rotation pattern. That is, we want to discover whether other alternatives might be preferable, and for this purpose we want to bring some of the alternatives into view in an explicit manner. If, for example, it were imperative that all the staffing goals for Type Duty 2 are to be attained, then the goal cost

for Type duty 2 should be set to some larger number, like 10. Such a rearrangement of goal deviation penalties results in a new rotation pattern as is summarized in Table A-8.

Type duty	1	2	3
Staffing Goal	35	125	50
Best Possible Schedule	<u>35</u>	<u>125</u>	<u>35</u>
Divergence	0	0	15

Table A-8 Goals vs Best Possible, Type Duty 2 Pre-emptive

Comparing the rotations summarized in Tables A-7 and A-8, we see that 25 persons (about 12%) had to be rotated to the second choice in the rotation sequence, and 10 persons (about 5%) had to be rotated late. The dollar cost, on the other hand, was reduced from \$570,000 to \$405,000 in going from Table A-7 to A-8, since the staffing goals for Type 2 personnel were met by holding back personnel due to be rotated.

To conclude this example, we show , in Table A-9, the complete solution for the case illustrated in Table A-7, with all the rotations and costs for each type of duty and length of service. Looking at the first row of the table, the FROM NODE is characterized by Type Duty and time on tour, as is the TO NODE. In other words, the first row of data is from Type Duty 1--tour length 1, to Type Duty 1--tour length 2. Since this is the preferred transition, the GOAL COST--i.e., the penalty for deviation from this goal--is 0; and, since no PCS move is involved

in remaining in Type Duty 1, DOLLAR COST is also \$0. The ARC FLOW is 10, indicating that all 10 persons starting in Type 1 Duty with less than 1 full period service (time on tour 1) were transferred. ARC COST(G), as explained earlier, is the total penalty, i.e. the product of GOAL COST and ARC FLOW. In this case, ARC COST is 0 since no penalty was incurred. ARC COST(\$) is the dollar equivalent of ARC COST(G), and is also \$0, since no PCS cost was incurred on this arc. Finally, the column labelled MARG COST(G) indicates the rate of increase of goal cost per unit increase of flow (of personnel) on an arc (for small increases). On our first row, the absence of any MARG COST indicates that this row is part of the solution, and no additional cost would be incurred by adding any personnel flow on this arc, i.e., bringing this arc into the solution, since it already is in the solution. The values of 6 which do appear in later rows indicate that the goal cost or penalty increase per additional person on such a rotation, should such a rotation be allowed as part of the solution, is 6. This, in turn, implies that a change of 6 units of goal costs would be necessary before this arc could be considered for inclusion in the final solution--i.e. before any personnel would be considered for this rotation.

FROM NODE	TO NODE	GOAL COST	DOLLAR COST	ARC FLOW	ARC COST(G)	ARC COST(\$)	MARG COST(G)
Type Duty	Time Tour	Type Duty	Time Tour				
1	1	1	2	0	0	0	
1	1	2	1	4	5000	0	
1	1	3	1	2	20000	0	
1	2	1	3	4	0	0	6
1	2	2	1	2	5000	10*	25000
1	2	3	1	0	20000	5	100000
1	3	1	3	4	0	0	0
1	3	2	1	2	5000	10	20*
1	3	3	1	0	20000	0	0
2	1	1	1	2	5000	0	0
2	1	2	2	0	0	40	0
2	1	3	1	4	9000	0	0
2	2	1	1	2	5000	0	0
2	2	2	3	0	0	40	0
2	2	3	1	4	9000	0	0
2	3	1	1	0	5000	25	0
2	3	2	3	4	0	0	0
2	3	3	1	2	9000	15	30*
3	1	1	1	4	20000	0	0
3	1	2	1	2	9000	0	0
3	1	3	2	0	0	15	0
3	2	1	2	4	20000	0	0
3	2	2	1	2	9000	0	0
3	2	3	3	0	0	15	0
3	3	1	1	2	20000	0	0
3	3	2	1	0	9000	15	0
3	3	3	3	4	0	0	0
Totals				195	60	570000	

* indicates a penalized rotation

Table A-9. Actual Flows in the Optimized Network

OPTIMIZATION AND MONITORING PROCEDURES

The operating procedures for GPSSR can be divided into two categories, depending upon whether an optimization function (using the analytical goal-programming network model) or a monitoring function is desired. The monitoring function can also be considered a sub-function of the optimization function as well as an (important) function in its own right. The procedures presented in the next two sections are those necessary for executing the GPSSR sequence of programs. As mentioned earlier, plans are in progress for more "intelligent" user-friendly procedures which would require minimal knowledge of computer-related concepts.

OPTIMIZATION PROCEDURES

During this developmental stage of GPSSR, with only a few representative DCs being considered, it is possible to maintain these DCs data on a disk file, and thus "on-line." The production version of GPSSR, however, must be able to access data for any of the entire set of DCs. For this reason, the data for all DCs will reside on magnetic tape. Thus, the initial step of GPSSR's execution must be one of reading the specified DCs data from magnetic tape onto disk, creating an "on-line" environment for that DC. The five modules of GPSSR must be executed in the order listed, thereby restricting the user to issue commands in a specified order, as shown below:

MODULE "Extraction of Data"

EXTRACT nn : Mounts tape of DC nn, and extracts relevant fields on to a disk file, producing an "on-line" environment for DC nn.

MODULE "Calculation of Transition Rates"

TRANSIT nn : Computes transition rates for DC nn.

GPSSMO nn : Smooths transition rates for DC nn.

MODULE "Network Generator"

GPSSRO nn : Performs the advanced start for the network optimizer of DC nn. Computes staffing levels based on historical and user defined rates (before flexibility introduced).

MODULE "Network Optimizer"

VICNET nn : Computes solution to sea/shore rotation problem for DC nn. i.e. minimizes deviations from goals while minimizing actual dollar costs.

MODULE "Report Generator"

REPORT1 nn : Produces a summary report for DC nn.

Note that if the user wishes to provide input in the form of smoothing or flexibility parameters or "overriding" transition rates, he or she must do so by editing (creating or modifying) a NAMELIST file prior to the execution of the appropriate task. The format of the NAMELIST file is as follows:

SNAME ALPHA=.1,
 DELTA=.1,

•
•
•
\$END

MONITORING PROCEDURES

A critical role of GPSSR is one of monitoring past and present consequences of personnel management. Using EMR data, the system can generate a variety of descriptive statistics and display them in formats which are meaningful for managers. This important monitoring function can be achieved by means of the three "plotting" tasks, directed by the commands PLOTSS, PLOTTT, and PLOTSM, as follows:

PLOTTSS nn : Plots personnel at sea versus personnel on shore for
DC nn.

PLOTTT nn : Plots time on tour for DC nn.

PLOTSM nn : Provides plots of smoothed data for DC nn.

In order to monitor the community, the optimizer need not be invoked, nor must the monitoring be used when optimizing; however, the extraction modules discussed in the previous section must be called before the plotting routines. The next section shows the required processing order of the GPSSR functions and modules.

APPENDIX B

EXPONENTIAL SMOOTHING

Given a series of historical rates, which may be trendy and noisy, a common method for estimating a "true" current rate is exponential smoothing. Proceeding in the manner of an exponential function, this technique weights current data more heavily than the earlier data. The user selects a parameter, α , which determines how much additional weight should be given to the current year. Choosing $\alpha = 0$ gives equal weight to all years. This is equivalent to taking the mean of the time series as the estimate for the current value. Conversely, choosing $\alpha = 1$ uses only the current year as the estimate of the "true" value for the series.

In our case the time series consists of historical rates of promotion, loss, etc., for each of the past 4 years. We need an estimate of the rates for the next 4 or 5 years. The rates for even a stable community tend to oscillate somewhat, and, when the oscillation is not too great, as is true for most stable communities, the mean rates would be most appropriate. For an expanding community, however, especially a community which has started expanding less than 4 years ago, the mean rates are less appropriate than a weighted series in which the most recent year is given greater weight.

The exponential smoothing formula is not generally given in closed form, but is usually given recursively. If $R(n)$ is the rate for year n ,

and $S(n-1)$ is the smoothed rate for year $n-1$, then the formula for $S(n)$ is given by

$$S(n) = \alpha * R(n) + (1 - \alpha) * S(n-1)$$

For the first year

$$S(1) = R(1)$$

In our program, we actually take a modified $S(n)$, $S'(n)$, where

$$S'(n) = \alpha * S(n) + (1 - \alpha) * M$$

where M is the mean rate for the entire series.¹

As an example, suppose the data for 5 years are 10,8,11,9,12. This represents a series 8.5,9,9.5,10,10.5 with "noise" of 1.5,-1,1.5,-1,1.5 "added". Choosing $\alpha = 0$ gives an estimate of $S = 10$, which is too low an estimate for the current average value of the series. Choosing $\alpha = 1$ gives an estimate of $S = 12$, which is too high. The next value of the series will actually be 10, but with a "true" value of 11. An α of .3 gives the estimate $S = 10.10$, while an α of .1 gives an estimate of 10.005. Both of these are a better estimate than the mean for the "true" value of the series.

¹ The user is allowed to vary α for each rate. There is no compelling unique choice of α , beyond the obvious observations that, if the community is stable, $\alpha = 0$ is the correct choice, while if data earlier than the current year is irrelevant, then $\alpha = 1$ should be used. Many textbooks recommend choosing α between .01 and .3. However, we have run a number of tests with different choices of α , but without any decisive results. The user not skilled in time series analysis is advised to use the default values for α .

REFERENCES

- [1] Charnes, A., and Cooper, W. W., **Management Models and Industrial Applications of Linear Programming**, Volumes I and II, New York, Wiley, 1961.
- [2] Charnes, A., Cooper, W. W., and Ferguson, R., "Optimal Estimation of Executive Compensation by Linear Programming," **Management Science**, Vol. I, No. 2, January, 1955, pp. 138-151.
- [3] Charnes, A., Cooper, W. W., and Niehaus, R. J., "Analytical Models for Manpower Decisions," **Naval Research Reviews**, May, 1972, pp. 1-9.
- [4] Lovegren, V., **Informatics, Analysis, and Solutions of a Class of Constrained Network Problems**, Dissertation, The University of Texas, 1983.
- [5] McLaughlin, Mark, "The Sea/Shore Rotation Model (SSRM) User's Manual," B-K Dynamics, Inc. Publication TR-3-497, Rockville, MD 20852 (1982).
- [6] Niehaus, R. **Studies in Manpower Planning** (Washington: U.S. Navy Office of Civilian Manpower Management, 1972).

USER'S MANUAL FOR THE MARK 1 GPSSR SYSTEM

The GPSSR system may logically be divided into six parts; the first five parts take data from the EMR plus inputs from the analyst to assist the policy-maker in developing an optimal rotation policy; the sixth part is the graphic interface, currently a SAS interface, which enables the policy-makers and detailers to monitor the current status of the policy, showing in as useful a manner as possible the current state of personnel on tour. The first five parts are as follows: the first is the extraction section, which reads historical data from the EMR; the second section calculates the smoothed rates of promotion, demotion, attrition, accession and, if desired, rotation from the historical data; the third section is the network generator; the fourth section is the network optimizer; and the fifth section is the report generator, which extracts from the network optimizer output whatever information the analyst needs.

We emphasize that the current system is a prototype, whose operation has been determined by responses to immediate demands, and which is set up to allow the greatest flexibility to respond to requests for data which is quantitatively and qualitatively different from what we formerly thought would be needed. A production system would have fewer of what are now known to be unneeded options. Normally, in operation we expect that the smoothed transition rates will have already been prepared, so the policy-maker should not have to worry about this part of the system. However, for completeness, we include sections on the operation of these subsystems.

1. NORMAL OPERATION

We assume that the data has already been extracted from the historical EMR and smoothed, and is in the data files:

MKxxx This is the file containing all the smoothed transition rates, i.e., the rates of promotion, demotion, etc., broken down by paygrade, length of service, time on tour, and type duty; The xxx is the Rate Code Number of the community or other identifier.

ACxxx This is a file of historical accessions to the community. An accession here means to the community, not to the Navy, so anyone transferring from another community would be counted as an accession. Accessions are broken down by length of service, time on tour, etc., just as the transition rates above are.

STxxxxyy This is the file which contains the starting inventory of community xxx in year yy.

R1xxx This file contains information on the historical rates of rotation, broken down by rank and length of service, and gives the rate at which personnel transfer from one type duty to another, e.g. the percentage of personnel in type duty 1 who rotate to type duties 2, 3, 4, 5, and 6 etc.; it does not include information about when personnel rotate. This is to allow the analyst to give tour lengths without specifying the type of rotation to be performed at the end of the tour, but allowing historical percentages to determine the type of duty to which personnel shall be rotated. As before, xxx is the Rate Code Number of the community.

R2xxx This is a file of historical rotation rates broken down by paygrade, length of service, type duty and time on tour.

Using R2xxx, the analyst can project the effects of a continuation of historical rates.

R1xxx and R2xxx are mainly useful for analysing how actual rotations differed from stated policies, and are not necessary for the normal operation of the system.

In the modelling effort, the network optimizer may be taken as a black box, and the modeller will spend most of the effort on the Network Generator. This program takes the various weights, goals and penalties and generates the network to be optimized. Most of the model's flexibility is accessed at this time, so the analyst has a number of data files which may be modified. In an effort to make the effort tractable, however, default files have been set up, and any modifications irrelevant to the immediate needs of a specific analysis may be skipped. Hence the analyst may restrict attention to just a few aspects of sea/shore rotation at a time. The files which may be modified are as follows:

ACOxxxx DATE A file which, after running GPSM000, contains the average (smoothed) number of actual accessions to the community, but which may be adjusted to reflect future authorizations of personnel. xxxx is the Rate Code Number of the Community being analysed. See Fig. 1. The numbers are by year and type duty. All these accessions are assumed to be E3, Length of Service 1 year, and Time on Tour 1 year.

FILE: ACO4000 DATA

YEAR	ACCESSIONS INTO TYPE DUTY 1-6					
84	8	768	0	133	8	0
85	8	768	0	133	8	0
86	8	768	0	133	8	0
87	8	768	0	133	8	0
88	8	768	0	133	8	0
89	8	768	0	133	8	0

FIGURE 1. USER DEFINED ACCESSIONS

CONSTS DATA A file with constants used by the code. The constants in this file should probably not be touched by the user, since most must agree with other values present in the program, and this data file is primarily to make program modifications easier, except for

NYEAR The number of periods on the model's horizon, i.e., the number of periods the model is to run; and

IFY The first year on the model's horizon; for example, if the last year of data was from 1983, then IFY would be 84 or 1984.

The other constants in the data file are:

MAXLOS Here, MAXLOS is the maximum length of service over which to disaggregate. For example, if MAXLOS = 10, all personnel with Length of Service greater than or equal to 10 will be lumped into a single category 10+. In addition to MAXLOS, which here must be less than or equal to 32 because of the pro-assigned lengths of certain data structures.

MAXPG The maximum paygrade to disaggregate. If, for example, MAXPG = 5, then all paygrades greater than 5 would be lumped

together into a single category 5+. Currently, MAXPG must be less than or equal to 9.

NTOUR The longest possible tour length to consider. As with the preceding values, if NTOUR = 6, all tour lengths greater than 6 will be lumped together. Currently, 6 is the longest possible tour that can be considered.

NTD The number of types of duty to be considered. The model is currently limited to 6 distinct types of duty, but will be extended to 8 duty types to allow for accompanied and unaccompanied overseas duty.

RTSET DATA A file containing data on a user specified rotation policy. See Fig. 2. The user may specify:

MAXROT The desired length of a tour, broken down by paygrade, type duty, and LOS class. For now, three LOS classes are considered: 1-4, 5-17, and 18+. The data is arranged as follows: in the first block are the values for duty type 1, first the rotation tour lengths for the seven paygrades 3 to 9, Length of Service category 1, then for the seven paygrades and Length of Service category 2, then tour length for the seven paygrades and Length of Service category 3; this pattern is then repeated for the remaining 5 types of duty.

ITOTD Preferred TO duty, the first number in the first row is the first choice for type duty 1; the second number is the first choice of rotation assignment for type duty 2, etc. The first number in the second row is the second choice of rotation assignment for type duty 1, the second number is the second choice for type duty 2, etc.

IRTCST Dollar Cost of a rotation. The first number is from duty 1
to duty 1; the second is from duty 2 to duty 1; etc.

FILE: RTSET DATA

&RTSETS

MAXROT =

2, 2, 2, 2, 2, 2, 2, 2, 2,
2, 2, 2, 2, 2, 3, 3, 3, 3,
3, 3, 3,

5, 5, 5, 5, 5, 5, 5, 5, 5,
5, 5, 5, 5, 5, 3, 3, 3, 3,
3, 3, 3,

5, 5, 5, 5, 5, 5, 5, 5, 5,
5, 5, 5, 5, 5, 3, 3, 3, 3,
3, 3, 3,

5, 5, 5, 5, 5, 5, 5, 5, 5,
5, 5, 5, 5, 5, 3, 3, 3, 3,
3, 3, 3,

5, 5, 5, 5, 5, 5, 5, 5, 5,
5, 5, 5, 5, 5, 3, 3, 3, 3,
3, 3, 3,

2, 2, 2, 2, 2, 2, 2, 2, 2,
2, 2, 2, 2, 2, 3, 3, 3, 3,
3, 3, 3,

ITOTD=

4, 1, 6, 5, 2, 3,
2, 6, 1, 2, 4, 4,
3, 4, 5, 1, 3, 2,
5, 3, 2, 6, 1, 5,
6, 5, 4, 3, 6, 1,
1, 2, 3, 4, 5, 6,

IRTCST=

0, 10, 10, 10, 10, 10,
10, 0, 10, 10, 10, 10,
10, 10, 0, 10, 10, 10,
10, 10, 10, 0, 10, 10,
10, 10, 10, 10, 0, 10,
10, 10, 10, 10, 10, 0

&END

FIGURE 2. ROTATION POLICY PARAMETERS

GOALxxxx DATA This is the file with the goals for detailing community xxxx. The goals are the desired (or billeted, or whatever the analyst thinks appropriate) staffing levels, by type duty and paygrade. The numbers are, first, type duty 1, paygrade 3; second type duty 1, paygrade 4; etc. See Fig. 3.

FILE: GOAL4000 DATA

GOALS	E3	E4	E5	E6	E7	E8	E9
	72	135	539	444	251	113	80
	2639	2224	1441	561	379	84	50
	1	2	11	6	7	2	3
	396	299	195	81	45	11	5
	27	22	19	16	23	15	9
	9	21	74	55	16	8	4

FIGURE 3. PERSONNEL STAFFING GOALS

PARAMS DATA These data may be entered interactively, or, alternatively, a file PARAMS DATA may be created with appropriate parameters in a NAMELIST format. See Fig. 4. The parameters which may be set are:

DELTA A flexibility parameter. Historical rates will be allowed to vary by no more than DELTA during a run. DELTA is currently expressed as a fraction; thus DELTA = .5 would allow a variance of no more than 50%.

GCOST An array of penalties for exceeding staffing goals. The numbers in the figure are as follows: The first 72 is the goal for E3 personnel in type duty 1; the second number, 135, is for E4 personnel in type duty 1, etc. In general, the

columns represent paygrades E3 to E9, while the rows represent the 6 duty types.

ICSTA1 Cost of accessions within the historical range specified by ACOxxxx and DELTA.

ICSTA2 Cost of accessions outside the historical range.

LCOST Per person cost of underachieving staffing goals.

ICSTR1 penalty costs to be assigned to rotations. Eventually, these costs should be determined automatically from the rotation policy; for now, however, they must be entered manually. A high cost should be assigned to improper rotations, and a low cost to preferred rotations. Thus, a high cost would be assigned to ICSTR1(1,1) since this represents keeping a person in type duty 1 beyond the stated rotation period.

ROTSET If TRUE, use stated rotation policy; if FALSE use historical rotation rates. Should normally be true.

ICOSTL Penalty assigned to losses. For now, this should probably be 0; it is provided as a variable for future extensions, when the model will have provisions for taking into account the cost of losses in determining an optimal rotation policy.

ICSTPD Penalty assigned to promotions and demotions. For now, this should be 0, but is included for future extensions of the model.

ACCSET If TRUE, use assigned accession rates; if FALSE, use historical rates.

FILE: PARAMS DATA

&PARAM

```

DELTA= 0.000000000E+00,
GCOST=      1,      1,      1,      1,      1,      1,
           1,      1,      1,      1,      1,      1,
           1,      1,      1,      1,      1,      1,
           1,      1,      1,      1,      1,      1,
           1,      1,      1,      1,      1,      1,
           1,      1,      1,      1,      1,      1,

ICSTA2= 5, LCOST= -5,
ICSTR1=    12,      0,      2,      4,      6,      8,
           0,      12,      6,      4,      2,      8,
           4,      6,      12,      2,      8,      0,
           0,      0,      0,      0,      0,      0,
           0,      0,      0,      0,      0,      0,
           0,      0,      0,      0,      0,      0,

ROTSET= T, ACCSET= T, ICSTL2= 5
&END

```

FIGURE 4. USER DEFINED PARAMETERS FOR GPSSR

A sample PARMS DATA file is currently provided; if, however, FIGPS EXEC is altered, the user may alter parameters interactively. In FIGPS EXEC, file 15 should be changed from PARAMS DATA to TERM if the interactive option is desired.

When the user is satisfied with the data in the above mentioned files, the network generator file is run by typing

FIGPS xxxx yy Where xxxx is the community Rate Code Number, and yy is the last year for which data is available, i.e., the first year for the model to begin generating its network. yy should agree with the yy in the STxxxxyy file.

GPSSR1 The actual network generator. The matrix generator produces a file.

ARCSxxx DATA Where xxxx is the detailing community; this file is then read into the network optimizer.

2. NETWORK OPTIMIZER

The network optimizer, VICNET, is a very sophisticated piece of code, but one which can be treated as a "black box". The user merely types:

FINET xxxx Where xxxx is the detailing community, and
VICNET Which is the actual program.

and the network optimizer automatically calculates the rotation policy that minimized both penalty costs and dollar costs--a substantial improvement over systems that only simulate a rotation, and leave the analyst to search for an optimal policy by trial and error or "stubby pencil" methods.

3. REPORT GENERATOR

A preliminary report generator has been developed to provide a summary of the rotations computed as optimal by the network optimizer. This is accessed (minimally) by calling:

FIREP xxxx yy Where xxxx is the name of the community, and yy is the
 last year of historical data (used for start
 inventories).
REPORT The program that actually produces the report.

Actually, for every distinct application, modifications to the report generator will probably be necessary. The current report only gives a summary of some of the possible information available to the analyst; consequently, since not all categories are present, the columns do not (and should not in general be expected to) add up. See Fig. 5.

PERIOD: FY85	TYPE DUTY: CONUS SHORE DUTY							
PAYGRADE	E3	E4	E5	E6	E7	E8	E9	TOTAL
STATED ROTATION POLICY	2	2	2	2	2	3	3	
INITIAL INVENTORY	75	301	643	321	209	78	63	1690
PROMOTIONS INTO PAYGRADE	0	36	171	162	0	0	0	369
LOSSES OUT OF PAYGRADE	0	0	0	0	0	0	0	0
ACCESSIONS INTO PAYGRADE	30	36	60	30	17	9	4	186
SCHEDULED INVENTORY	72	135	539	444	251	113	80	1634
PAYGRADE GOALS	72	135	539	444	251	113	80	1634
DEVIATIONS FROM GOALS	0	0	0	9	0	0	0	0

PERIOD: FY85	TYPE DUTY: ARDUOUS SEA DUTY							
PAYGRADE	E3	E4	E5	E6	E7	E8	E9	TOTAL
STATED ROTATION POLICY	5	5	5	5	5	3	3	
INITIAL INVENTORY	2857	2429	1256	650	293	60	45	7590
PROMOTIONS INTO PAYGRADE	0	0	0	0	0	0	0	0
LOSSES OUT OF PAYGRADE	0	0	0	0	0	0	0	0
ACCESSIONS INTO PAYGRADE	1274	653	159	50	35	7	4	2182
SCHEDULED INVENTORY	2639	2224	1441	561	379	84	50	7378
PAYGRADE GOALS	2639	2224	1441	561	379	84	50	7378
DEVIATIONS FROM GOALS	0	0	0	0	0	0	0	0

PERIOD: FY85	TYPE DUTY: OVERSEA SHORE DUTY							
PAYGRADE	E3	E4	E5	E6	E7	E8	E9	TOTAL
STATED ROTATION POLICY	5	5	5	5	5	3	3	
INITIAL INVENTORY	12	24	55	39	12	6	4	152
PROMOTIONS INTO PAYGRADE	0	0	0	0	0	0	0	0
LOSSES OUT OF PAYGRADE	0	0	0	0	0	0	0	0
ACCESSIONS INTO PAYGRADE	1	0	1	0	0	0	0	2
SCHEDULED INVENTORY	1	2	11	6	7	2	3	32
PAYGRADE GOALS	1	2	11	6	7	2	3	32
DEVIATIONS FROM GOALS	0	0	0	0	0	0	0	0

PERIOD: FY85	TYPE DUTY: NON-ROTATED SEA DUTY							
PAYGRADE	E3	E4	E5	E6	E7	E8	E9	TOTAL
STATED ROTATION POLICY	5	5	5	5	5	3	3	
INITIAL INVENTORY	570	490	505	544	255	81	58	2503
PROMOTIONS INTO PAYGRADE	0	7	3	0	0	0	0	10
LOSSES OUT OF PAYGRADE	0	0	0	0	0	0	0	0
ACCESSIONS INTO PAYGRADE	197	90	17	0	5	0	0	309

SCHEDULED INVENTORY	395	299	195	81	45	11	5	1032
PAYGRADE GOALS	395	299	195	81	45	11	5	1032
DEVIATIONS FROM GOALS	0	0	0	0	0	0	0	0

PERIOD: FY85	TYPE DUTY: NEUTRAL DUTY							
PAYGRADE	E3	E4	E5	E6	E7	E8	E9	TOTAL
STATED ROTATION POLICY	5	5	5	5	5	3	3	
INITIAL INVENTORY	23	32	64	24	28	13	3	187
PROMOTIONS INTO PAYGRADE	0	0	0	0	0	0	0	0
LOSSES OUT OF PAYGRADE	0	0	0	0	0	0	0	0
ACCESSIONS INTO PAYGRADE	15	6	0	1	1	0	0	23
SCHEDULED INVENTORY	27	22	19	16	23	15	9	131
PAYGRADE GOALS	27	22	19	16	23	15	9	131
DEVIATIONS FROM GOALS	0	0	0	0	0	0	0	0

PERIOD: FY85	TYPE DUTY: PREFERRED OVERSEAS SHORE DUTY							
PAYGRADE	E3	E4	E5	E6	E7	E8	E9	TOTAL
STATED ROTATION POLICY	2	2	2	2	2	3	3	
INITIAL INVENTORY	4	10	23	14	7	2	3	63
PROMOTIONS INTO PAYGRADE	0	0	0	0	0	0	0	0
LOSSES OUT OF PAYGRADE	0	0	0	0	0	0	0	0
ACCESSIONS INTO PAYGRADE	2	7	5	1	0	0	0	15
SCHEDULED INVENTORY	9	21	74	55	16	8	4	187
PAYGRADE GOALS	9	21	74	55	16	8	4	187
DEVIATIONS FROM GOALS	0	0	0	0	0	0	0	0

PERIOD: FY85	TYPE DUTY: TOTAL							
PAYGRADE	E3	E4	E5	E6	E7	E8	E9	TOTAL
INITIAL INVENTORY	3541	3286	2546	1592	804	240	176	12185
PROMOTIONS INTO PAYGRADE	0	43	174	162	0	0	0	379
LOSSES OUT OF PAYGRADE	0	0	0	0	0	0	0	0
ACCESSION INTO PAYGRADE	1519	792	242	82	58	16	8	2717
SCHEDULED INVENTORY	3144	2703	2279	1163	721	233	151	10394
PAYGRADE GOALS	3144	2703	2279	1163	721	233	151	10394
DEVIATIONS FROM GOALS	0	0	0	0	0	0	0	0

FIGURE 5. SAMPLE REPORT OUTPUT.

For the sake of completeness, operation of the modules for extracting the data and producing the smoothed transition rates is described below.

4. EXTRACTION

The commands under the CMS operating system are (currently):

FIEXT yy x Where yy is the year when the data were taken, and x is the disk on which the EMR resides. Note that copies of the historical EMR must be loaded onto disk from tape before FIEXT can be called. FIEXT, as well as other commands beginning with FI... are EXECs to initialize CMS, before running the actual programs.

EXTRACT8 Program to extract the data from the file FYyy (where yy is the year when the historical EMR was snap-shotted) and put the data into FYyyEXT.

FIEXT2 yy xx Where yy is the year to be extracted, and xx is the Rate Code Number of the community to be examined.

DREDUCE This is the program that computes from the current year (supplied by the user) and the years on the EMR the length of service and time on tour, and which extracts from the EMR a single detailing community for analysis. The program will prompt the user for the year when the data were taken, since this is not kept on the historical EMR tapes, and for Rate Code Number xx supplied above. At some point, we hope to automate this process. The resulting files are called FYyyxx, yy and xx as described above for FIEXT2.

The above programs cannot be run from the Navy account on the UT system, at this time, since the account does not have enough disk space to

accommodate simultaneously all communities. If needed, this situation will be altered; however, we anticipate that the extracted transitions will be kept permanently on disk, so these extractions will not be needed for normal day to day operation of the model.

5. CALCULATING SMOOTHED RATES

The smoothing programs can be run from the Navy account, but this should not normally be necessary. The files FYxxxxyy DATA must have already been extracted from the historical EMR, where xxx is the detailing community's Rate Code Number, and yy is the year when the EMR was generated.

The smoothing programs may then be run in the following order:

TRANSIT This program counts all the transitions, presenting the totals of all personnel promoted, demoted, etc., as well as the totals of all personnel, broken down by length of service, time on tour, etc.

The original historical snapshots of the EMR each contain data for several communities from a single year. In order to calculate the smoothed rates, the system must do the bookkeeping chore of concatenating the separate yearly files into a 5 year amalgam of data for a single community. This is done by

FISEVSAS yy xx Where yy is the first year in the series and xx is the last year in the series.

SEVSAS This is the program that actually concatenates the annual files, while writing out additional files which will eventually be available for SAS graphics.

At this point, the analyst may adjust the data files to be used by the next program, GPSMOO, which turns the transition totals into smoothed transition rates. However, since this section was primarily intended to be used by the programmers rather than analysts, some of the data values cannot be changed without altering the actual code, or inconsistencies will result. This will only be true during the prototype stage of development; the production version of the code will not suffer from these difficulties. The relevant files are:

SPARMS DATA A file containing relevant parameters in NAMELIST format, similar to PARMS, described in Section 1 are:

ALPHA The smoothing constant mentioned in the GPSSE report, Appendix B, by type duty and paygrade, thus the first entry is for type duty 1, paygrade 3, the second entry is for duty 1, paygrade 4, etc.

ROTSET If TRUE, use user defined rotation policy; if FALSE, use historical rates.

SCONST DATA A file containing relevant constants in NAMELIST format, similar to CONSTS, described in Section 1. The constants are:

MAXLOS The number of distinct categories of length of service; currently, MAXLOS = 3, where category 1 represents lengths of service from 1 to 4,; category 2 represents lengths of service from 5 to 17; and category 3 represents all lengths of service of 18 years or more. The value of 3 is necessary to maintain compatibility with parts of the program, and is for the convenience of the programmer; it is not to be set by the user.

MAXPG The maximum paygrade to disaggregate. If, for example, MAXPG = 5, then all paygrades greater than 5 would be lumped together into a single category 5+.

FRSTYR The first year of historical data.

STOPYR The last year of historical data.

NTOUR The longest possible tour length to consider. As with the precedign values, if NTOUR = 6, all tour lengths greater than 6 will be limped together.

NTD The number of types of duty to be considered.

SRTSET DATA A file containing data on a user specified rotation policy, similar to RTSETS, described in Section 1. The user may specify:

MAXROT The desired length of a tour, broken down by paygrade and type duty. The first value, as above, is for type duty 1, paygrade 3, length of service category 1 (i.e., 1 to 4 years of service). The remaingin values are as in RTSETS.

ITOTD Preferred TO duty, the first number is the first choice for type duty 1; the second number is the second choice of rotation assignment for type duty 1, etc.

IRTCST Dollar Cost of a rotation. The first number is from duty 1 to duty 1; the second is from duty 1 to duty 2; etc.

When the analyst is satisfied with these values, the smoothing routine is called:

GP3M000 xx This command automatically calls the program that smooths the transition rates. xx is the Rate Code Number of the community being analysed.

6. GRAPHIC MONITORING FACILITY

A number of SAS procedures have been provided for the analyst.

Currently, these are designed to produce a binary "meta" plot file, which may be disposed to a four-color zeta pen plotter; alternatively, the files may be modified to produce results at the user's terminal, but the quality of SAS interactive graphics on a non-graphic terminal is not very good. Only a quick look is possible before disposing the results to the pen plotter. Plots produced by the Navy analyst can be available within 48 hours after the plot request is submitted, however.

END
FILMED

5-86

DTIC